

# Scalable Bayesian Inference for the Inverse Temperature of a Hidden Potts Model

Matthew T. Moores      Anthony N. Pettitt  
Kerrie Mengersen

March 30, 2015

## Abstract

The inverse temperature parameter of the Potts model governs the strength of spatial cohesion and therefore has a substantial influence over the resulting model fit. A difficulty arises from the dependence of an intractable normalising constant on the value of this parameter and thus there is no closed form solution for sampling from the posterior distribution directly. There are a variety of computational approaches for sampling from the posterior without evaluating the normalising constant. These algorithms differ in their levels of accuracy and their scalability for datasets of realistic size. This study compares pseudolikelihood, the exchange algorithm, path sampling, and approximate Bayesian computation. We assess their accuracy and scalability using synthetic data as well as 2D remote sensing and 3D medical images. Our findings provide guidance on selecting a suitable algorithm for Bayesian image analysis. For nontrivial images, this necessarily involves some degree of approximation to produce an acceptable compromise between accuracy and computational cost.

*Keywords:* Approximate Bayesian computation, Composite likelihood, Hidden Markov random field, Image analysis, Pseudo-marginal method, Thermodynamic integration.

---

Matthew T. Moores is Research Associate, Department of Statistics, University of Warwick, Coventry CV4 7AL, UK (E-mail: [M.T.Moores@warwick.ac.uk](mailto:M.T.Moores@warwick.ac.uk)). Anthony N. Pettitt is Professor and Kerrie Mengersen is Professor, Mathematical Sciences School, Queensland University of Technology, Brisbane, Queensland 4001, Australia (E-mail addresses: [a.pettitt@qut.edu.au](mailto:a.pettitt@qut.edu.au) and [k.mengersen@qut.edu.au](mailto:k.mengersen@qut.edu.au)). The authors thank the Radiation Oncology Mater Centre, Queensland Health, for providing the CT scans. Landsat imagery courtesy of NASA Goddard Space Flight Center and U.S. Geological Survey. Computational resources and services used in this work were provided by the HPC and Research Support Group, Queensland University of Technology.

# 1 Introduction

Markov random field (MRF) models have seen widespread use in image analysis since their introduction by [Besag \[1974\]](#), as surveyed by [Winkler \[2003\]](#) and [Li \[2009\]](#). A MRF is a generalisation of the Markovian dependence structure to more than one dimension: satellite imagery has two spatial dimensions, while medical images such as computed tomography (CT) are three-dimensional. The hidden [Potts \[1952\]](#) model employs a latent MRF on discrete states to describe spatial dependence between adjacent neighbours. The degree of dependence in the model is governed by a parameter known as the inverse temperature due to its origin in statistical physics. It is difficult to set this parameter by trial and error, particularly for noisy images. Rather than using a fixed value, it would be preferable to estimate the inverse temperature as part of the model. However, the intractable normalising constant of the Potts model depends on the value of the inverse temperature, which means that there is no closed form solution for estimating its posterior distribution.

[Rydén and Titterington \[1998\]](#) derived a pseudolikelihood (PL) approximation [\[Besag, 1975\]](#) to the intractable posterior density. [Gelman and Meng \[1998\]](#) instead approximated the ratio of normalising constants using thermodynamic integration (TI), also known as path sampling. [Møller et al. \[2006\]](#) introduced an auxiliary variable method that gives an exact MCMC algorithm for the special case of a 2 component Potts model, also known as an Ising model. [Murray et al. \[2006\]](#) proposed a variant of the exact method known as the exchange algorithm or multiple auxiliary variable method (MAVM). By replacing the expensive perfect sampling step [\[Propp and Wilson, 1996\]](#) with an approximation such as Gibbs sampling, [Cucala et al. \[2009\]](#) developed an approximate exchange algorithm that can be applied for Potts models with  $k > 2$ . [Friel et al. \[2009\]](#) introduced the reduced dependence approximation (RDA) which uses recursion to calculate the normalising constant on small sub-lattices. [McGrory et al. \[2012\]](#) generalised RDA to an irregular lattice. [Grelaud et al. \[2009\]](#) used the sufficient statistic of the Potts model to estimate the inverse temperature using approximate Bayesian computation (ABC). [Everitt \[2012\]](#) combined the approximate exchange algorithm with particle Markov chain Monte Carlo (PMCMC) and also implemented ABC with sequential Monte Carlo (SMC-ABC) for the Ising model. Although all of these methods work very well in theory, the largest image used in any of these papers to demonstrate an algorithm is less than ten thousand pixels. This does not give a reliable indication of how these algorithms might perform when applied to images of a more substantial size.

Images containing multiple megapixels are now commonplace, from digital photography and high-definition (HD) video to medical imaging and remote sensing. Table 1 illustrates how the number of pixels translates to real world scale for various types of images, including the area covered by a satellite image, the number of axial slices of a CT scan, and the number of frames of HD video. Multispectral images from the Landsat 7 [\[NASA, 2011\]](#) and Landsat 8 [\[USGS, 2014\]](#) satellites have a spatial resolution of 30 metres per pixel (area of  $900\text{m}^2$ ). A Landsat image covers an area of approximately 170km north-

Number of pixels	Satellite (900m <sup>2</sup> /px)	CT slices (512×512)	HD Video 1080p
$2^6$	0.06km <sup>2</sup>	...	...
$5^6$	14.06km <sup>2</sup>	0.1	...
$10^6$	900.00km <sup>2</sup>	3.8	0.5
$15^6$	10251.56km <sup>2</sup>	43.5	5.5

Table 1: Scale of common types of images.

south by 183km east-west. Tomographic reconstructions such as CT scans are usually represented as 3D image stacks, with  $512 \times 512$  pixels per slice. The pixel resolution and slice width varies depending on the clinical protocol. A single frame of HD video is typically  $1920 \times 1080$  pixels with 24 frames per second (fps), although higher frame rates and resolutions (such as ultra high definition, UHD) are available. The volumes of data involved in video processing necessitate specialised methods, which are beyond the scope of this paper. For examples of Bayesian methods for video analysis, see [Simoncelli \[1999\]](#), [Minvielle et al. \[2010\]](#), and the references therein. The remainder of the discussion will focus on static images.

This is the first study to systematically investigate the scalability of methods for intractable likelihoods. In this paper we compare the performance of pseudo-likelihood (PL), the approximate exchange algorithm (MAVM), thermodynamic integration (TI), and approximate Bayesian computation (ABC). By conducting these experiments we are able to address questions such as how the computational cost increases with the size of the images, as well as how much accuracy is lost by using faster, more approximate methods. The imaging datasets, which include 2D remote sensing and 3D medical imaging, are described in Section 2. We define the hidden Potts model in Section 3 and examine the properties of its intractable normalising constant. The various computational methods for estimating the inverse temperature are reviewed in Section 4. We present our experimental results in Section 5 and conclude the article with a discussion.

## 2 Example datasets

### 2.1 Synthetic data

The inverse temperature cannot be directly observed, so the only way to measure the accuracy of a method is to use simulated data. We have adapted a method that was proposed by [Cook et al. \[2006\]](#) for the purpose of assessing the correctness of software for fitting a Bayesian model. Their central concept is to draw values of the model parameters from the prior and simulate data using those parameters, then fit the model and compare the estimate of the posterior distribution to the known value. We do not make use of the  $\chi^2$  test devised by [Cook et al.](#), since our goal is not to evaluate the hypothesis that our

implementation of the algorithm is correct. We know that all four algorithms are approximations, therefore our aim is to compare the performance of these algorithms under various conditions.

We use a uniform prior to draw values of the inverse temperature. This prior is adopted because it reflects current practice and enables comparison between our results and other recent studies. We use informative, conjugate priors for the means and variances of the mixture components in our simulation study. This avoids the problem of label switching and corresponds with the model that we use in our applications to real data. We use a fixed number of 3 mixture components to generate 6 datasets containing 20 images each. The images within a dataset all have the same dimensions. The small images have  $2^6$  pixels ( $8 \times 8$  for 2D images, or  $4 \times 4 \times 4$  for 3D). The medium images have  $5^6$  ( $125 \times 125$ , or  $25 \times 25 \times 25$ ). The large images have  $10^6$  ( $1000 \times 1000$ , or  $100 \times 100 \times 100$ ). This enables us to study how the performance of the algorithms changes as the size of the images increases, as well as how well these algorithms generalise to a 3D lattice.

## 2.2 Satellite remote sensing

Landsat imagery of the Greater Brisbane region was converted to the Normalised Difference Vegetation Index (NDVI). This index is calculated using the visible red (*RED*) and near-infrared (*NIR*) bands:

$$\text{NDVI} = \frac{\text{NIR} - \text{RED}}{\text{NIR} + \text{RED}} \quad (1)$$

NDVI is generally used to detect vegetated areas, with an expected value for vegetation being between 0.3 and 0.8. In this analysis, we are interested in detecting clusters which relate to vegetated areas. The size of the data set does not enable us to individually view the satellite image of Brisbane, however, as indicated in Figure 1b, a mixture model may be an appropriate means of determining clusters, and attributes of clusters, specifically those related to vegetation. In addition, a further goal of data analysis is detecting loss (or gain) of moderate sized vegetative areas through time. A mixture model may also suit this task, as pixel movement between clusters may indicate a change in habitat or land use.

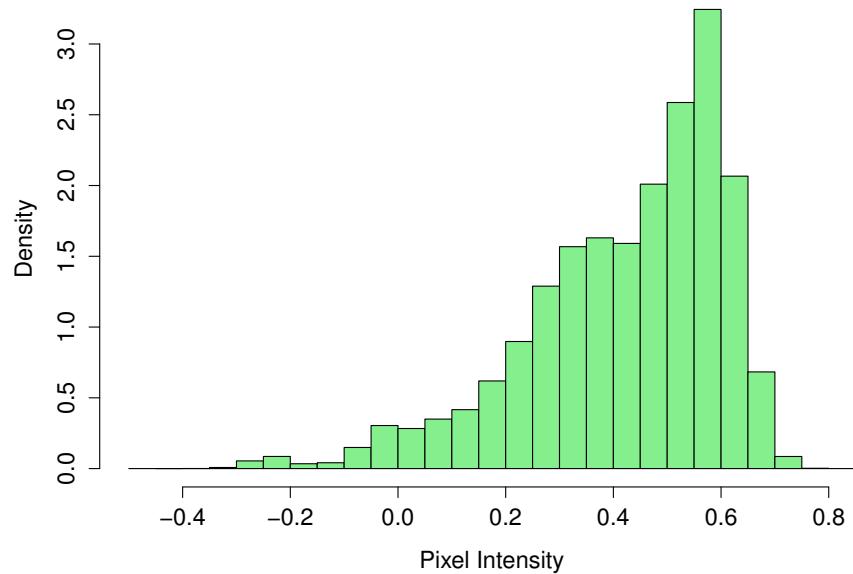
## 2.3 Medical images

Computed tomography (CT) is a technology for reconstructing a three-dimensional image by filtered back-projection from X-ray beams [Kalender, 2011]. The pixel intensities in a CT scan are linearly related to the electron density (ED) within the object that is imaged. CT scanners are calibrated in Hounsfield units (HU) with water at 0 HU and air at -1024 HU.

Rather than using biological images, which are difficult for the untrained eye to interpret, in this paper we have used 27 CT scans of an ED phantom (CIRS, Inc. model 062). This object is manufactured from epoxy and contains

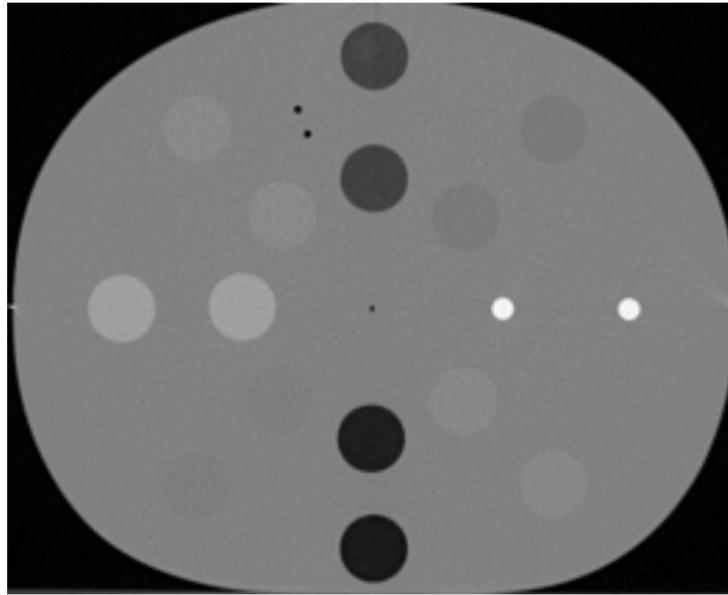


(a) Satellite image of Brisbane, Australia.

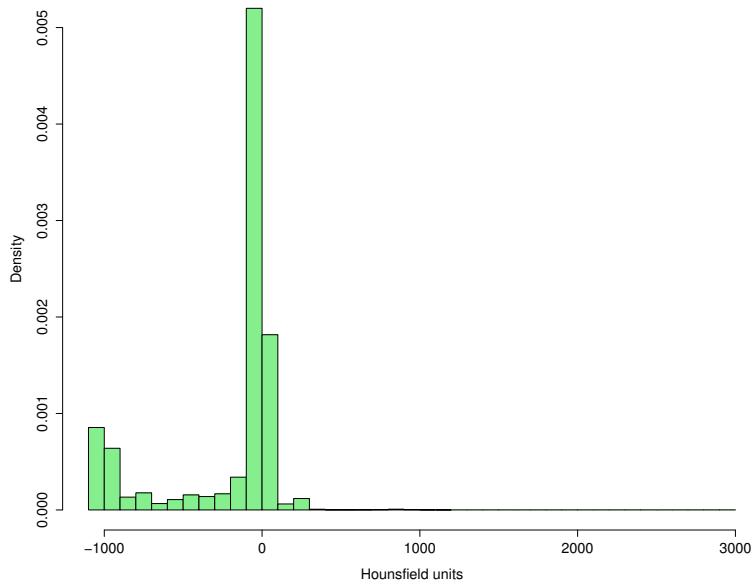


(b) Histogram of NDVI values.

Figure 1: A satellite image and the distribution of pixel intensities.



(a) CT scan of the ED phantom.



(b) Histogram of CT numbers (HU).

Figure 2: An axial slice of a CT scan and the distribution of pixel intensities.

cylindrical inserts that mimic the X-ray absorption of human tissue. Since the geometry and density of this object is known, the CT scans can be used to explore the effect of inverse temperature on segmentation accuracy. An example slice of a CT scan of the ED phantom is illustrated in Figure 2. The cylindrical inserts are arranged in pairs, clockwise from the top: lung (exhale); adipose (fatty tissue); dense (cortical) bone; muscle; lung (inhale); breast (50% fat); spongy (trabecular) bone; and liver. The body of the phantom is composed of water-equivalent solid and therefore should have a mean pixel intensity of 0 HU.

### 3 Hidden Potts model

Image segmentation can be viewed as the task of labelling the observed pixels  $\mathbf{y}$  according to a finite set of discrete states  $\mathbf{z} \in \{1, \dots, k\}$ . The hidden Potts model allows for spatial correlation between neighbouring labels in the form of a Markov random field. The latent labels follow a Gibbs distribution, which is specified in terms of its conditional probabilities:

$$p(z_i | z_{\setminus i}, \beta) = \frac{\exp \{ \beta \sum_{i \sim \ell} \delta(z_i, z_{\ell}) \}}{\sum_{j=1}^k \exp \{ \beta \sum_{i \sim \ell} \delta(j, z_{\ell}) \}} \quad (2)$$

where  $\beta$  is the inverse temperature,  $z_{\setminus i}$  represents all of the labels except  $z_i$ ,  $i \sim \ell$  are the neighbouring pixels of  $i$ , and  $\delta(u, v)$  is the Kronecker delta function. Thus,  $\sum_{i \sim \ell} \delta(z_i, z_{\ell})$  is a count of the neighbours that share the same label.

If the labels  $z_i$  are indexed row-wise, the nearest (first-order) neighbours  $i \sim \ell$  in a regular 2D lattice with  $c$  columns are  $\{z_{i-1}, z_{i-c}, z_{i+c}, z_{i+1}\}$ . Pixels situated at the boundary of the image domain have less than four neighbours. Likewise, voxels in a regular 3D lattice have a maximum of 6 first-order neighbours. These neighbourhood relationships are reciprocal, so  $h \in i \sim \ell$  implies  $i \in h \sim \ell$ . If  $\mathcal{E}$  is the set of all unique neighbour pairs, or edges in the image lattice, then  $|\mathcal{E}| = 2(n - \sqrt{n})$  for a square lattice and  $3(n - n^{2/3})$  for a cube.

The observation equation links the latent labels to the corresponding pixel values:

$$p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) = \prod_{i=1}^n p(y_i | z_i, \theta_{z_i}) \quad (3)$$

where  $\theta_j$  are the parameters that govern the distribution of the pixel values with label  $j$ . The hidden Potts model can thus be viewed as a spatially-correlated generalisation of the finite mixture model [Rydén and Titterington, 1998]. Green and Richardson [2002] used a Poisson likelihood for (3), with intensity  $\lambda_j$ . Instead we follow Geman and Geman [1984], Alston et al. [2007], and many others in assuming that the pixels with label  $j$  share a common mean  $\mu_j$  corrupted by additive Gaussian noise with variance  $\sigma_j^2$ :

$$y_i | z_i = j, \mu_j, \sigma_j^2 \sim \mathcal{N}(\mu_j, \sigma_j^2) \quad (4)$$

The Gibbs distribution is a member of the exponential family and so there is a sufficient statistic for this model, as noted by [Grelaud et al. \[2009\]](#):

$$S(\mathbf{z}) = \sum_{i \sim \ell \in \mathcal{E}} \delta(z_i, z_\ell) \quad (5)$$

This statistic represents the total number of like neighbour pairs in the image. The likelihood  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}, \beta)$  can therefore be factorised into  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})p(S(\mathbf{z}) | \beta)$ , where the second factor does not depend on the observed data, but only on the sufficient statistic. The joint posterior is then:

$$p(\boldsymbol{\theta}, \beta, \mathbf{z} | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})\pi(\boldsymbol{\theta})p(S(\mathbf{z}) | \beta)\pi(\beta) \quad (6)$$

The conditional distributions  $p(\boldsymbol{\theta} | \mathbf{z}, \mathbf{y})$  and  $p(z_i | z_{\setminus i}, \beta, y_i, \boldsymbol{\theta}_{z_i})$  can be simulated using Gibbs sampling, but  $p(\beta | \mathbf{y}, \mathbf{z}, \boldsymbol{\theta})$  involves an intractable normalising constant  $\mathcal{C}(\beta)$ :

$$p(\beta | \mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) \propto p(S(\mathbf{z}) | \beta)\pi(\beta) \quad (7)$$

$$\propto \frac{\exp\{\beta S(\mathbf{z})\}}{\mathcal{C}(\beta)}\pi(\beta) \quad (8)$$

The normalising constant is also known as a partition function in statistical physics. It has computational complexity of  $\mathcal{O}(nk^n)$ , since it involves a sum over all possible combinations of the labels  $\mathbf{z} \in \mathcal{Z}$ :

$$\mathcal{C}(\beta) = \sum_{\mathbf{z} \in \mathcal{Z}} \exp\{\beta S(\mathbf{z})\} \quad (9)$$

It is infeasible to calculate this value exactly for nontrivial images, thus computational approximations are required.

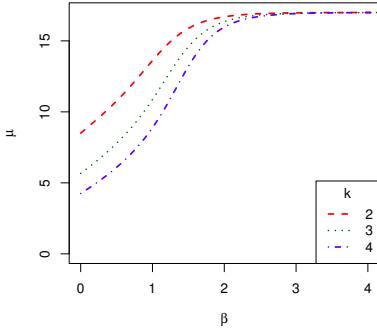
The conditional expectation of  $S(\mathbf{z})$  given  $\beta$  can be expressed in terms of the normalising constant:

$$\mathbb{E}_{\mathbf{z} | \beta}[S(\mathbf{z})] = \frac{d}{d\beta} \log\{\mathcal{C}(\beta)\} \quad (10)$$

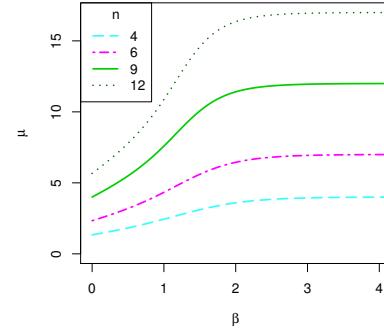
As  $\beta$  approaches infinity, all of the pixels in the image are almost surely assigned the same label, thus the expectation of  $S(\mathbf{z})$  approaches the total number of edges  $|\mathcal{E}|$  asymptotically, while the variance approaches zero. When  $\beta = 0$ , (2) simplifies to  $\left(\sum_j \exp\{0\}\right)^{-1}$ , hence the probability of any pair of neighbours being assigned the same label follows an independent Bernoulli distribution with  $p = k^{-1}$ . In this case,  $S(\mathbf{z})$  follows a Binomial distribution with expectation  $|\mathcal{E}|/k$  and variance  $|\mathcal{E}|k^{-1}(1 - k^{-1})$ . In a finite image lattice the distribution of  $S(\mathbf{z})$  changes smoothly between these two extremes, as illustrated by Figure 3, but its computation is intractable for nontrivial images.

The Potts model undergoes a phase transition at the critical value of  $\beta$ , switching from a disordered to an ordered state. [Potts \[1952\]](#) showed that the critical value for a 2D lattice on a cylinder can be calculated exactly according to:

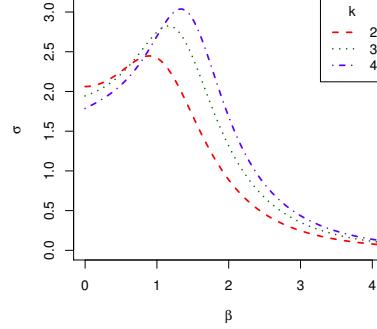
$$\beta_{crit} = \log\left\{1 + \sqrt{k}\right\} \quad (11)$$



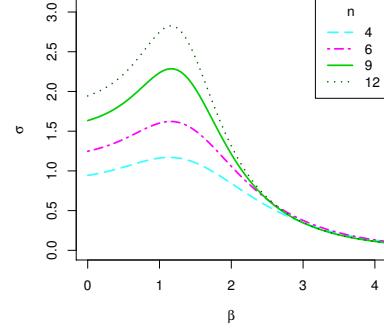
(a) Expectation for  $n = 12$  and increasing values of  $k \in \{2, 3, 4\}$ .



(b) Expectation for  $k = 3$  and increasing values of  $n \in \{4, 6, 9, 12\}$ .



(c) Standard deviation for  $n = 12$  and increasing values of  $k \in \{2, 3, 4\}$ .



(d) Standard deviation for  $k = 3$  and increasing values of  $n \in \{4, 6, 9, 12\}$ .

Figure 3: Distribution of the sufficient statistic of the Potts model for increasing values of the inverse temperature  $\beta$ , the number of pixels  $n$ , and the number of unique labels  $k$ . The expectation and standard deviation of  $S(\mathbf{z})$  were calculated exactly, using a brute force method.

The periodic boundary condition does not apply to the lattices considered in this paper, so for example the critical value for the images in Figure 3 is different to (11). However, the error introduced by the finite boundary diminishes as  $n$  increases. Figure 4 shows that (11) is very accurate in predicting the behaviour of  $S(\mathbf{z})$  for a 2D image with a maximum value of  $S(\mathbf{z})$  for first-order neighbours of 1,954,672. This corresponds to the satellite image described in Section 2.2, with  $k=6$  mixture components and  $\beta_{crit} \approx 1.24$ .  $S(\mathbf{z})$  is approximated by simulation using the algorithm of Swendsen and Wang [1987]. Figure 4a shows that the gradient of the expectation becomes very steep near the critical region, which is reflected in the super-exponential increase of the standard deviation illustrated by Figure 4c. As  $n \rightarrow \infty$  the derivative of the likelihood at the critical point, and hence the variance, is unbounded [Pickard, 1987]. This heteroskedasticity has important implications for many of the methods discussed in Section 4.

There is no exact formula for  $\beta_{crit}$  in 3D images, although Hajduković [1983] developed an empirical approximation that works reasonably well:

$$\beta_{crit} \approx \frac{2}{3} \log \left\{ \frac{1}{2} \left( \sqrt{2} + \sqrt{4k-2} \right) \right\} \quad (12)$$

The behaviour of  $\mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})]$  for a 3D image with  $k = 9$  and  $\beta_{crit} \approx 0.86$  is illustrated by Figure 4b. This is typical of the CT scans described in Section 2.3, with  $|\mathcal{E}|$  approximately equal to three million. The standard deviation is shown in Figure 4d. It is evident that (12) has underestimated  $\beta_{crit}$  since the maximum value of the standard deviation occurs at  $\beta = 0.9$ .

## 4 Computational methods

In this section we describe the four algorithms that are evaluated in Section 5. These methods provide alternative means to simulate parameter values from (8) without computing the intractable normalising constant. We describe the algorithms in terms of Markov chain Monte Carlo (MCMC) to enable direct comparison, although we also mention other approaches where applicable, such as particle-based (SMC and PMCMC) methods. Reference implementations of all of these methods are available from various sources described below, but for the purpose of comparison we have reimplemented the algorithms using RcppArmadillo [Eddelbuettel and Sanderson, 2014]. An R source package containing our code is available in Moores and Mengersen [2015].

### 4.1 Pseudolikelihood and composite likelihood

Pseudolikelihood is the simplest of the methods that we have considered and also the fastest. Rydén and Titterington [1998] showed that the intractable distribution (8) could be approximated using the product of the conditional densities given by (2). This enables updates for the inverse temperature at iteration  $t$  to be simulated using a Metropolis-Hastings (M-H) step, as shown in Algorithm 1.

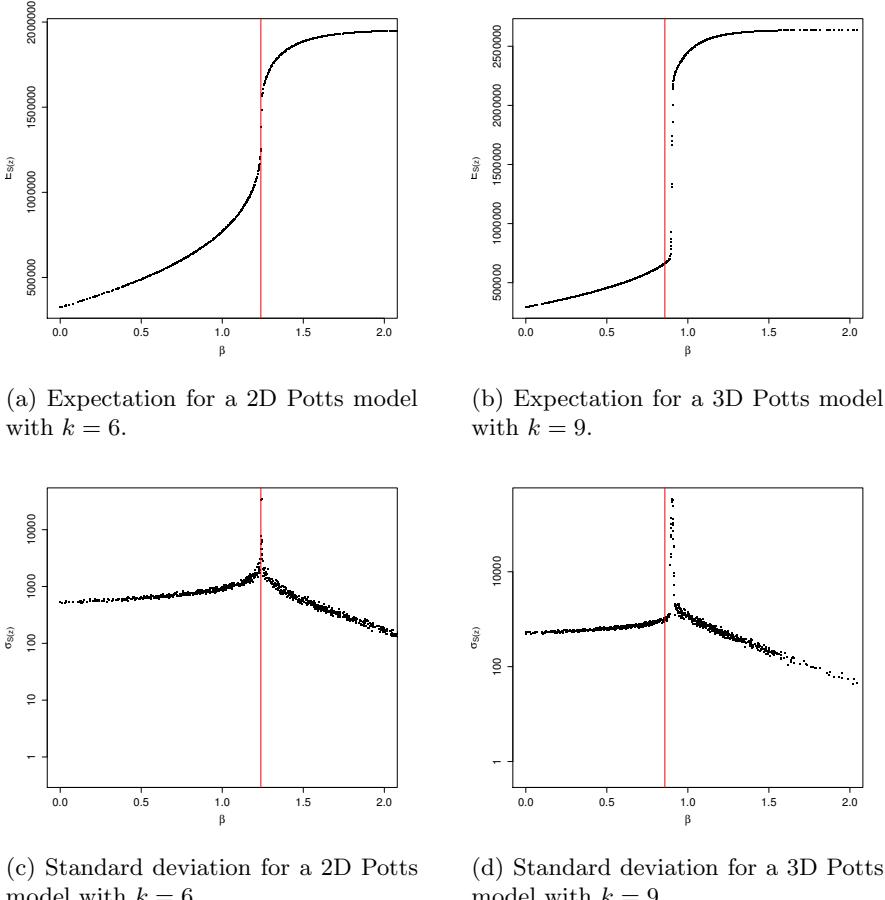


Figure 4: Approximate distribution of  $S(\mathbf{z})$  using Swendsen-Wang for a 2D image with  $k = 6$  and  $|\mathcal{E}| \approx 2 \times 10^6$  (left) and a 3D image with  $k = 9$  and  $|\mathcal{E}| \approx 3 \times 10^6$  (right). The vertical line is the critical value of  $\beta$ .

---

**Algorithm 1** MCMC with pseudolikelihood (PL)

---

```

1: Initialise  $\beta_0, \mu_0, \sigma_0^2, \mathbf{z}_0$ 
2: for  $t \in 1 \dots T$  do
3:   Update the labels  $z_i \sim p(y_i|z_i, \mu_{z_i}, \sigma_{z_i}^2) p(z_i|z_{\setminus i}, \beta) \forall i \in \{1, \dots, n\}$ 
4:   Calculate sufficient statistics  $\bar{y}_j, s_j^2 \forall z_i = j, \forall j \in \{1, \dots, k\}$ 
5:   Update the noise parameters  $\mu_j, \sigma_j^2 \sim p(\bar{y}_j, s_j^2|\mu_j, \sigma_j^2) \pi(\mu_j|\sigma_j^2) \pi(\sigma_j^2)$ 
6:   Draw proposed parameter value  $\beta' \sim q(\beta'|\beta_{t-1})$ 
7:   Approximate  $p(\mathbf{z}|\beta')$  and  $p(\mathbf{z}|\beta_{t-1})$  using:

```

$$p(\mathbf{z}|\beta) \approx \prod_{i=1}^n p(z_i|z_{\setminus i}, \beta) \quad (13)$$

```

8:   Calculate the M-H ratio  $\rho = \min \left( 1, \frac{p(\mathbf{z}|\beta')\pi(\beta')q(\beta_{t-1}|\beta')}{p(\mathbf{z}|\beta_{t-1})\pi(\beta_{t-1})q(\beta'|\beta_{t-1})} \right)$ 
9:   Draw  $u \sim \text{Uniform}[0, 1]$ 
10:  if  $u < \rho$  then
11:     $\beta_t \leftarrow \beta'$ 
12:  else
13:     $\beta_t \leftarrow \beta_{t-1}$ 
14:  end if
15: end for

```

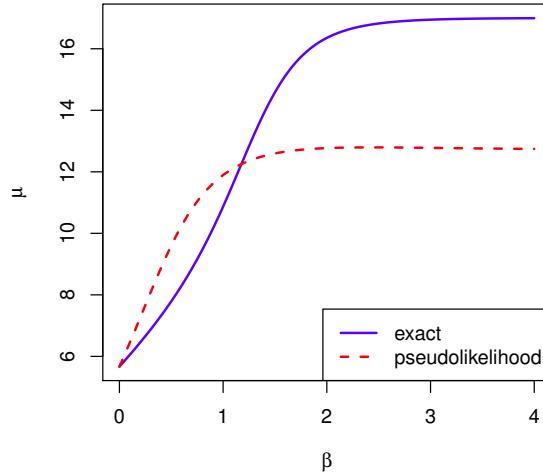
---

The M-H proposal density  $q(\beta'|\beta_{t-1})$  can be any distribution such that  $\int q(\beta'|\beta_{t-1}) d\beta' = 1$ . However, there is a tradeoff between exploring the full state space and ensuring that the probability of acceptance (at line 10) is sufficiently high. We use the adaptive random walk (RWMH) algorithm of [Garthwaite et al. \[2010\]](#), which automatically tunes the bandwidth of the proposal density to target a given M-H acceptance rate. When a symmetric proposal density is used,  $q(\beta'|\beta_{t-1}) = q(\beta_{t-1}|\beta')$  and so this term cancels out in the M-H ratio on line 8. The natural logarithm of  $\rho$  is used in practice to improve the numerical stability of Algorithm 1.

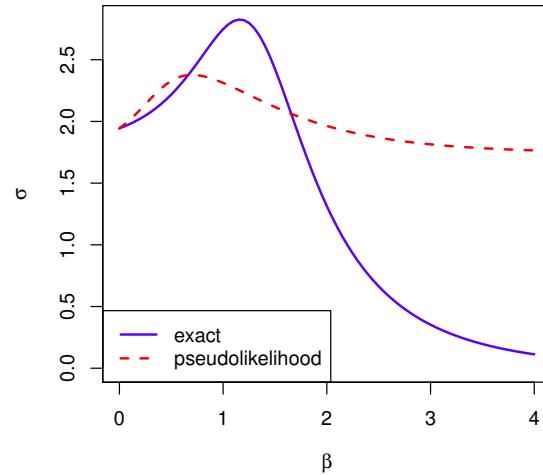
Pseudolikelihood is exact when  $\beta = 0$  and provides a reasonable approximation for small values of the inverse temperature. However, the approximation error increases rapidly for  $\beta \geq \beta_{crit}$ , as illustrated by Figure 5. This is due to long-range dependence between the labels, which is inadequately modelled by the local approximation. The implications of this inaccuracy for posterior inference will be demonstrated in Section 5.

[Rydén and Titterington \[1998\]](#) referred to Equation (13) in Algorithm 1 as point pseudolikelihood, since the conditional distributions are computed for each pixel individually. They suggested that the accuracy could be improved using block pseudolikelihood. This is where the likelihood is calculated exactly for small blocks of pixels, then (13) is modified to be the product of the blocks:

$$p(\mathbf{z}|\beta) \approx \prod_{i=1}^{N_B} p(\mathbf{z}_{B_i}|\mathbf{z}_{\setminus B_i}, \beta) \quad (14)$$



(a) Expectation.



(b) Standard deviation.

Figure 5: Approximation error of pseudolikelihood for  $n = 12, k = 3$  in comparison to the exact likelihood calculated using a brute force method: (a)  $\sum_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{z}) p(S(\mathbf{z})|\beta)$  using either Equation (8) or (13); (b)  $\sqrt{\sum_{\mathbf{z} \in \mathcal{Z}} (S(\mathbf{z}) - \mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})])^2 p(S(\mathbf{z})|\beta)}$

where  $N_B$  is the number of blocks,  $\mathbf{z}_{B_i}$  are the labels of the pixels in block  $B_i$ , and  $\mathbf{z}_{\setminus B_i}$  are all of the labels except for  $\mathbf{z}_{B_i}$ . This is a form of composite likelihood, where the likelihood function is approximated as a product of simplified factors [Varin et al., 2011]. Friel [2012] compared point pseudolikelihood to composite likelihood with blocks of  $3 \times 3$ ,  $4 \times 4$ ,  $5 \times 5$ , and  $6 \times 6$  pixels. Friel showed that (14) outperformed (13) for the Ising ( $k = 2$ ) model with  $\beta < \beta_{crit}$ . Okabayashi et al. [2011] discuss composite likelihood for the Potts model with  $k > 2$  and have provided an open source implementation in the R package **potts** [Geyer and Johnson, 2014].

Evaluating the conditional likelihood in (14) involves the normalising constant for  $\mathbf{z}_{B_i}$ , which is a sum over all of the possible configurations  $\mathcal{Z}_{B_i}$ . This is a limiting factor on the size of blocks that can be used. The brute force method that was used to compute Figure 3 and 5 is too computationally intensive for this purpose. Pettitt et al. [2003] showed that the normalising constant can be calculated exactly for a cylindrical lattice by computing eigenvalues of a  $k^r \times k^r$  matrix, where  $r$  is the smaller of the number of rows or columns. The value of (9) for a free boundary lattice can then be approximated using path sampling. Friel and Pettitt [2004] extended this method to larger lattices using a composite likelihood approach.

The reduced dependence approximation (RDA) is another form of composite likelihood. Reeves and Pettitt [2004] introduced a recursive algorithm to calculate the normalising constant using a lag- $r$  representation. Friel et al. [2009] divided the image lattice into sub-lattices of size  $r_1 < r$ , then approximated the normalising constant of the full lattice using RDA:

$$\mathcal{C}(\beta) \approx \frac{\mathcal{C}_{r_1 \times n}(\beta)^{r-r_1+1}}{\mathcal{C}_{r_1-1 \times n}(\beta)^{r-r_1}} \quad (15)$$

McGrory et al. [2009] compared RDA to pseudolikelihood and the exact method of Møller et al. [2006], reporting similar computational cost to pseudolikelihood but with improved accuracy in estimating  $\beta$ . Source code for RDA is available in the online supplementary material for McGrory et al. [2012].

## 4.2 Thermodynamic integration

Gelman and Meng [1998] derived an approximation to the log ratio of normalising constants using the path sampling identity:

$$\log \left\{ \frac{\mathcal{C}(\beta_{t-1})}{\mathcal{C}(\beta')} \right\} = \int_{\beta'}^{\beta_{t-1}} \mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})] d\beta \quad (16)$$

which follows from the definition of  $\mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$  in (10). The value of this definite integral can be approximated by simulating from the Gibbs distribution for fixed values of  $\beta$  and then interpolating between them. We used the Swendsen-Wang algorithm for simulating from  $\mathbf{z}|\beta$ , as mentioned in section 3. Figure 6a illustrates linear interpolation of  $\mathbf{S}(\mathbf{z})$  on a 2D lattice for  $k = 6$  and  $\beta$  ranging from 0 to 2 in increments of 0.05. This approximation can be precomputed

using the same method that was used to produce Figure 4. Figure 6b illustrates interpolation of  $S(\mathbf{z})$  on a 3D lattice for  $k = 9$ .

---

**Algorithm 2** Path sampling (TI)

---

```

1: Initialise  $\beta_0, \mu_0, \sigma_0^2, \mathbf{z}_0$ 
2: for  $t \in 1 \dots T$  do
3:   Update the labels  $z_i \sim p(y_i|z_i, \mu_{z_i}, \sigma_{z_i}^2) p(z_i|z_{\setminus i}, \beta) \forall i \in \{1, \dots, n\}$ 
4:   Calculate sufficient statistics  $S(\mathbf{z})$  and  $\bar{y}_j, s_j^2 \forall z_i = j, \forall j \in \{1, \dots, k\}$ 
5:   Update the noise parameters  $\mu_j, \sigma_j^2 \sim p(\bar{y}_j, s_j^2 | \mu_j, \sigma_j^2) \pi(\mu_j | \sigma_j^2) \pi(\sigma_j^2)$ 
6:   Draw proposed parameter value  $\beta' \sim q(\beta' | \beta_{t-1})$ 
7:   Estimate  $\mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})]$  for  $\beta \in \{\beta', \beta_{t-1}\}$  by interpolation
8:   Evaluate the definite integral in Equation (16)
9:   Calculate the log M-H acceptance ratio:

$$\log\{\rho\} = \min \left( 0, \log \left\{ \frac{\mathcal{C}(\beta_{t-1})}{\mathcal{C}(\beta')} \right\} + (\beta' - \beta_{t-1})S(\mathbf{z}) + \log \left\{ \frac{\pi(\beta')q(\beta_{t-1}|\beta')}{\pi(\beta_{t-1})q(\beta'|\beta_{t-1})} \right\} \right) \quad (17)$$

10:  Draw  $u \sim \text{Uniform}[0, 1]$ 
11:  if  $u < \rho$  then
12:     $\beta_t \leftarrow \beta'$ 
13:  else
14:     $\beta_t \leftarrow \beta_{t-1}$ 
15:  end if
16: end for

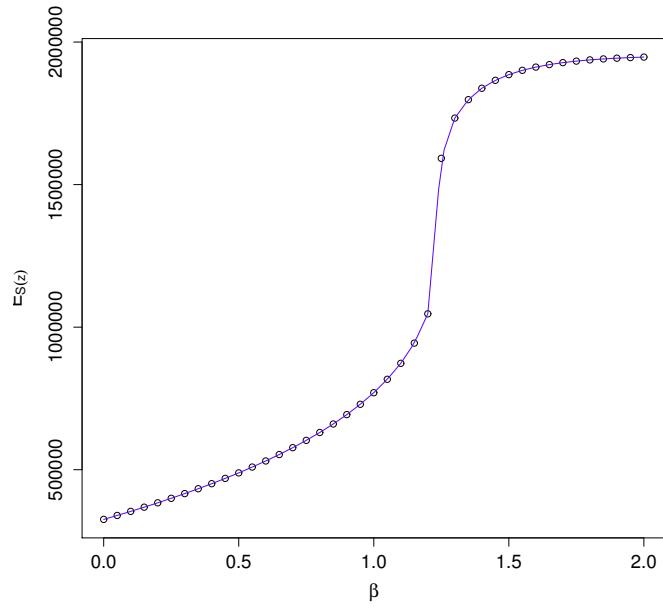
```

---

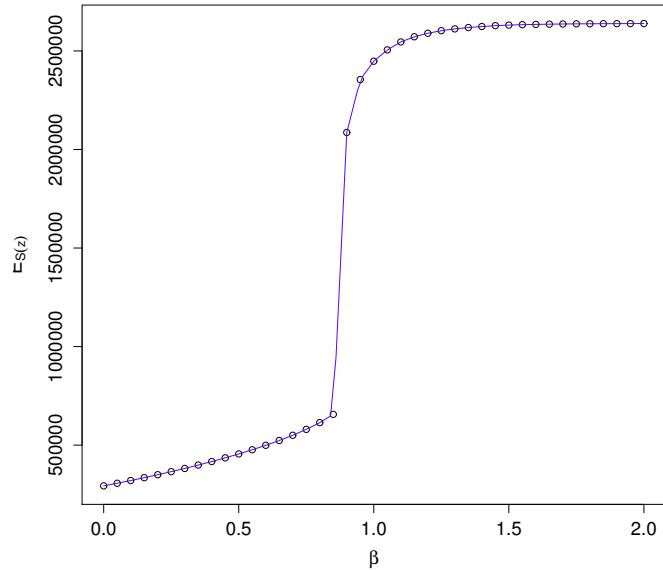
Path sampling is explained in further detail by [Chen et al. \[2000\]](#), chap. 5. A reference implementation in R is available from the website accompanying [Marin and Robert \[2007\]](#). This algorithm has an advantage over auxiliary variable methods such as the exchange algorithm or ABC because the additional simulations are performed prior to fitting the model, rather than at each iteration. This is particularly the case when analysing multiple images that all have approximately the same dimensions, as in the examples of section 5. However, the computational cost is still slightly higher than pseudolikelihood, which does not require a pre-computation step.

### 4.3 Pseudo-marginal methods

[Møller et al. \[2006\]](#) demonstrated that it is possible to simulate from the posterior distribution of  $\beta$  by introducing an auxiliary variable, so that  $\mathcal{C}(\beta)$  cancels out in the M-H ratio. The exchange algorithm of [Murray et al. \[2006\]](#) improves the performance of this method and avoids the need for a fixed estimate of  $\beta$ . The drawback of this approach is that it requires perfect sampling from the stationary distribution of the Potts model. This is possible using coupling from the past [[Propp and Wilson, 1996](#), [Huber, 2003](#)] but can be computationally prohibitive, particularly for large images. [McGrory et al. \[2009\]](#) reported that the time required for perfect sampling increased sharply for larger values of  $\beta$ .



(a) 2D Potts model with  $k = 6$ .



(b) 3D Potts model with  $k = 9$ .

Figure 6: Approximation of  $\mathbb{E}_{\mathbf{z}|\beta}[S(\mathbf{z})]$  by simulation for fixed values of  $\beta$ , with linear interpolation.

For this reason, [Cucala et al. \[2009\]](#) substitute 500 iterations of Gibbs sampling on the auxiliary variable to produce an approximate sample from its stationary distribution. The details of this method are shown in Algorithm 3.

---

**Algorithm 3** Approximate exchange algorithm (MAVM)

---

```

1: Initialise  $\beta_0, \mu_0, \sigma_0^2, \mathbf{z}_0$ 
2: for  $t \in 1 \dots T$  do
3:   Update the labels  $z_i \sim p(y_i|z_i, \mu_{z_i}, \sigma_{z_i}^2) p(z_i|z_{\setminus i}, \beta) \forall i \in \{1, \dots, n\}$ 
4:   Calculate sufficient statistics  $\mathbf{S}(\mathbf{z})$  and  $\bar{y}_j, s_j^2 \forall z_i = j, \forall j \in \{1, \dots, k\}$ 
5:   Update the noise parameters  $\mu_j, \sigma_j^2 \sim p(\bar{y}_j, s_j^2|\mu_j, \sigma_j^2) \pi(\mu_j|\sigma_j^2) \pi(\sigma_j^2)$ 
6:   Draw proposed parameter value  $\beta' \sim q(\beta'|\beta_{t-1})$ 
7:   Generate  $\mathbf{w}|\beta'$  by sampling from Equation (2)
8:   Calculate the M-H acceptance ratio according to Equation (8):

$$\rho = \min \left( 1, \frac{q(\beta_{t-1}|\beta') \pi(\beta') \exp \{ \beta' \mathbf{S}(\mathbf{z}) \} \mathcal{C}(\beta_{t-1})}{q(\beta'|\beta_{t-1}) \pi(\beta_{t-1}) \exp \{ \beta_{t-1} \mathbf{S}(\mathbf{z}) \} \mathcal{C}(\beta')} \frac{\exp \{ \beta_{t-1} \mathbf{S}(\mathbf{w}) \} \mathcal{C}(\beta')}{\exp \{ \beta' \mathbf{S}(\mathbf{w}) \} \mathcal{C}(\beta_{t-1})} \right) \quad (18)$$

9:   Draw  $u \sim \text{Uniform}[0, 1]$ 
10:  if  $u < \rho$  then
11:     $\beta_t \leftarrow \beta'$ 
12:  else
13:     $\beta_t \leftarrow \beta_{t-1}$ 
14:  end if
15: end for

```

---

The M-H ratio given by Equation (18) can be considerably simplified if a uniform prior is used for  $\beta$  and  $\beta'$  is drawn from a symmetric M-H proposal density. In this case,  $\log\{\rho\}$  simplifies to  $(\beta' - \beta_{t-1})\mathbf{S}(\mathbf{z}) + (\beta_{t-1} - \beta')\mathbf{S}(\mathbf{w})$ . The similarity with Equation (17) shows that the exchange algorithm is closely related to path sampling, since it is a form of importance sampling. An implementation of the exchange algorithm is available in the online supplementary material accompanying [Friel and Pettitt \[2011\]](#). [Everitt \[2012\]](#) provides source code for the approximate exchange algorithm with particle MCMC.

#### 4.4 Approximate Bayesian computation

Like the exchange algorithm, ABC uses an auxiliary variable  $\mathbf{w}$  to decide whether to accept or reject the proposed value of  $\beta'$ . Instead of a M-H ratio such as (18), the summary statistics of the auxiliary variable and the observed data are directly compared. The proposal is accepted if the distance between these summary statistics is within the ABC tolerance,  $\epsilon$ . This produces the following approximation when the labels  $\mathbf{z}$  are observed without error:

$$p(\beta | \mathbf{z}) \approx \pi_\epsilon(\beta | \|\mathbf{S}(\mathbf{w}) - \mathbf{S}(\mathbf{z})\| < \epsilon) \quad (19)$$

where  $\|\cdot\|$  is a suitable norm. In this paper we simply use the absolute difference between  $\mathbf{S}(\mathbf{w})$  and  $\mathbf{S}(\mathbf{z})$ , since the summary statistic (5) is univariate.  $\mathbf{S}(\mathbf{z})$  is

sufficient for  $\beta$ , as noted by [Grelaud et al. \[2009\]](#), therefore the ABC approximation (19) approaches the true posterior as  $n \rightarrow \infty$  and  $\epsilon \rightarrow 0$ . In practice there is a tradeoff between the number of parameter values that are accepted and the size of the ABC tolerance.

[Everitt \[2012\]](#) introduced ABC for noisy data, such as the Ising model, but only when the observations are discrete and so  $S(\mathbf{y})$  is defined. [Moores et al. \[2015\]](#) showed that ABC can also be applied when the domain of the observed data is continuous, such as with the additive Gaussian noise of Equation (4). In this type of model the posterior distribution of  $\beta$  does not depend directly on the observed data  $\mathbf{y}$ , as shown by Equation (7). The ABC approximation (19) for the hidden Potts model becomes:

$$p(\beta | \mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) \approx \pi_\epsilon(\beta | \|S(\mathbf{w}) - S(\mathbf{z})\| < \epsilon) \quad (20)$$

The other parameters can be updated using the current value of  $\beta$  and then the summary statistic  $S(\mathbf{z})$  can be computed from the current values of the labels, using an ABC within Gibbs approach as shown in Algorithm 4. Equation (20) can be considered as a moving target, since  $S(\mathbf{z})$  could change whenever the labels are updated. As a consequence, ABC with noisy data can be more prone to getting stuck in low probability regions of the parameter space.

[Grelaud et al. \[2009\]](#) used a rejection sampler, where the proposals are drawn independently from the prior and thus  $q(\beta' | \beta_{t-1}) = \pi(\beta)$ . Under a sparse or uninformative prior, such as the uniform prior used in this study, too many proposals are rejected for this approach to be viable. Instead we have based Algorithm 4 on the ABC-MCMC algorithm of [Marjoram et al. \[2003\]](#). This form of ABC algorithm is best suited for direct comparison with the other intractable likelihood methods in this article.

---

**Algorithm 4** ABC-MCMC with noisy data

---

```

1: Initialise  $\beta_0, \boldsymbol{\mu}_0, \boldsymbol{\sigma}_0^2, \mathbf{z}_0$ 
2: for  $t \in 1 \dots T$  do
3:   Update the labels  $z_i \sim p(y_i | z_i, \mu_{z_i}, \sigma_{z_i}^2) p(z_i | z_{\setminus i}, \beta) \forall i \in \{1, \dots, n\}$ 
4:   Calculate sufficient statistics  $S(\mathbf{z})$  and  $\bar{y}_j, s_j^2 \forall z_i = j, \forall j \in \{1, \dots, k\}$ 
5:   Update the noise parameters  $\mu_j, \sigma_j^2 \sim p(\bar{y}_j, s_j^2 | \mu_j, \sigma_j^2) \pi(\mu_j | \sigma_j^2) \pi(\sigma_j^2)$ 
6:   Draw proposed parameter value  $\beta' \sim q(\beta' | \beta_{t-1})$ 
7:   Generate  $\mathbf{w} | \beta'$  by sampling from Equation (2)
8:   Draw  $u \sim \text{Uniform}[0, 1]$ 
9:   if  $u < \frac{\pi(\beta') q(\beta_{t-1} | \beta')}{\pi(\beta_{t-1}) q(\beta' | \beta_{t-1})}$  and  $\|S(\mathbf{w}) - S(\mathbf{z})\| < \epsilon$  then
10:     $\beta_t \leftarrow \beta'$ 
11:   else
12:     $\beta_t \leftarrow \beta_{t-1}$ 
13:   end if
14: end for

```

---

There have been many extensions to ABC, as reviewed by [Marin et al. \[2012\]](#). Of particular interest are adaptive ABC algorithms that automatically

adjust the tolerance  $\epsilon$  and the proposal bandwidth  $\sigma_{MH}^2$ . ABC with sequential Monte Carlo (SMC-ABC) algorithms use a sequence of target distributions  $\pi_{\epsilon_t}(\beta \mid \|S(\mathbf{w}) - S(\mathbf{z})\| < \epsilon_t)$  such that  $\epsilon_1 > \epsilon_2 > \dots > \epsilon_T$ , where the number of SMC iterations  $T$  can be determined dynamically using a stopping rule. The SMC-ABC algorithm of [Drovandi and Pettitt \[2011\]](#) uses multiple MCMC steps for each SMC iteration, while the algorithm of [Del Moral et al. \[2012\]](#) uses multiple replicates of the summary statistics for each particle. [Everitt \[2012\]](#) has provided a MATLAB implementation of SMC-ABC with the online supplementary material accompanying his paper.

The computational efficiency of ABC is dominated by the cost of drawing updates to the auxiliary variable, as reported by [Everitt \[2012\]](#). Thus, we would expect that the execution time for ABC would be similar to the approximate exchange algorithm. Various approaches to improving this runtime have recently been proposed, such as the Gaussian process emulation of [Wilkinson \[2014\]](#), the “lazy ABC” of [Prangle \[2014\]](#), and methods involving auxiliary models [[Cabras et al., 2014](#), [Moores et al., 2015](#), [Buzbas and Rosenberg, 2015](#)].

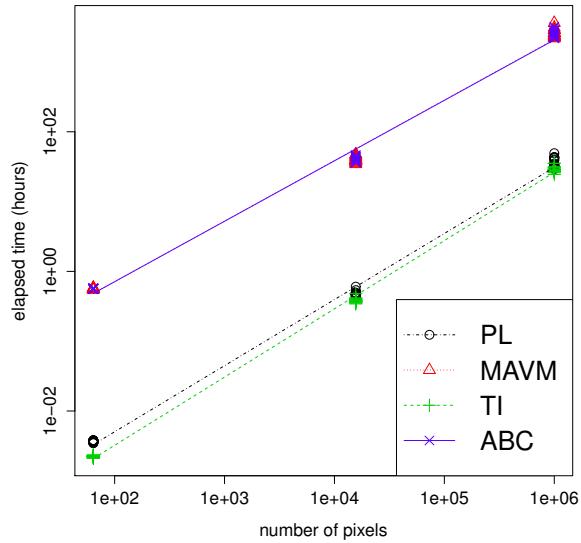
## 5 Experimental results

### 5.1 Simulation study

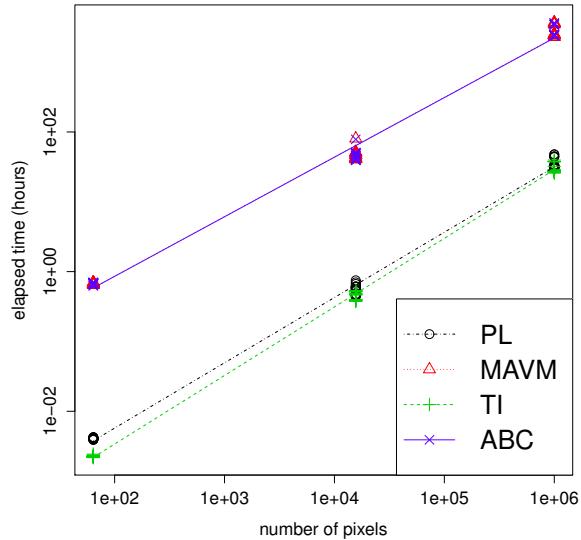
We generated the synthetic data using the R package `PottsUtils` [[Feng, 2008](#), [Feng and Tierney, 2014](#)] as explained in Section 2.1. We used informative priors  $\pi(\mu_1) \sim \mathcal{N}(-0.15, 0.05^2)$ ,  $\pi(\mu_2) \sim \mathcal{N}(0.05, 0.05^2)$ ,  $\pi(\mu_3) \sim \mathcal{N}(0.25, 0.05^2)$ , and  $\pi(\sigma_j^2) \sim \mathcal{IG}(1.5, 0.01)$  for the  $j \in 1 \dots 3$  mixture components and a uniform prior on the interval  $[0, 1.3\beta_{crit}]$  for the inverse temperature. For 2D images,  $\beta_{crit}$  was calculated using Equation (11) to be equal to 1.005. We used Equation (12) to obtain an approximate value for  $\beta_{crit}$  in 3D images of 0.552.

All four algorithms were run for 10,000 iterations on each image, discarding the first 5,000 as burn-in. The differences in elapsed runtime are illustrated by Figure 7. All of the algorithms scaled linearly with increasing numbers of pixels, but the approximate exchange algorithm (MAVM) and ABC were roughly two orders of magnitude slower than either pseudolikelihood (PL) or path sampling (TI). For 2D images with  $n = 10^6$  pixels the average runtime was 43 hours for MAVM, 42.5 hours for ABC, 36 minutes for PL and 29 minutes for TI. It took 30 minutes for precomputation of 64 values of  $\mathbb{E}_{\mathbf{z} \mid \beta}[S(\mathbf{z})]$  for path sampling, using 16 parallel threads. For 3D images with  $n = 10^6$  voxels the averages were 45 hours for MAVM or ABC, 35 minutes for PL and 30 minutes for TI. The precomputation step took 24 minutes for 44 values of  $\beta$ . Even allowing for the cost of precomputation, TI was still much faster than MAVM or ABC.

The distributions of the differences between the MCMC samples of  $\beta$  and the true value of the inverse temperature are illustrated by Figure 8 for 2D images with  $n = 10^6$ . Results for the other simulation studies are available in the supplementary material. It is evident that all four algorithms had difficulty with one of the images that had a very small value of  $\beta = 0.032$ . This appears to



(a) 2D images,  $k = 3$ .



(b) 3D images,  $k = 3$ .

Figure 7: Relationship between elapsed runtime and image size for pseudolikelihood (PL), the approximate exchange algorithm (MAVM), path sampling (TI), and approximate Bayesian computation (ABC). Both axes are on a logarithmic scale.

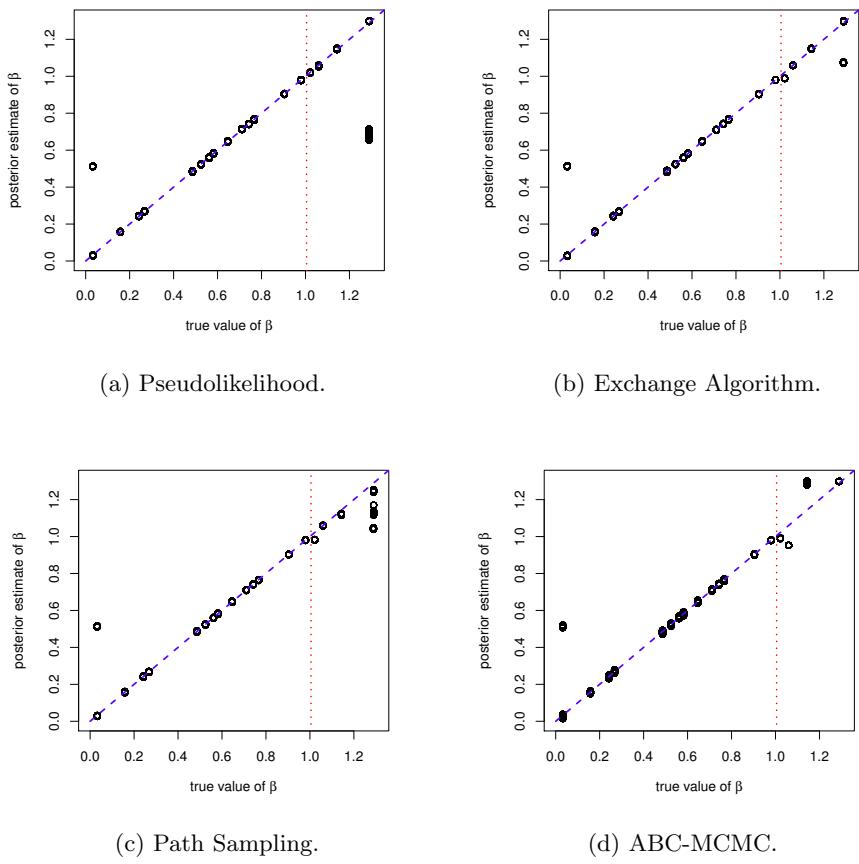


Figure 8: MCMC sampling error for  $\beta$  in simulated data for 20 images with 2 dimensions,  $k = 3$  and  $n = 10^6$ . The diagonal, dashed line indicates the true value of  $\beta$ , while the vertical, dotted line marks the critical point.

Method	95% HPD interval for $\beta$	Iterations	Elapsed	CPU time
Pseudolikelihood	[2.260; 2.309]	5K+55K	2.0h	15.5h
Path sampling	[1.248; 1.250]	5K+55K	1.4h	10.4h
Exchange algorithm	[1.275; 1.278]	5K+5K	50.8h	399.5h
ABC-MCMC	[1.316; 1.323]	5K+5K	50.6h	398.2h

Table 2: Results for the satellite image of Brisbane, Australia.

be due to the level of noise in the image. The algorithms produced an accurate estimate for another image with  $\beta = 0.033$  and much lower levels of noise. We did not observe this behaviour with any of the 3D images, nor with the smaller 2D images. We did, however, observe the errors increasing for values of  $\beta$  above the critical point, indicated by the vertical line in Figure 8. In the worst case, PL underestimated  $\beta$  by 0.61, with a 95% highest posterior density (HPD) interval of [0.653; 0.706] when the true value of  $\beta$  was 1.289. MAVM underestimated  $\beta$  for the same image by 0.215 and TI was out by 0.245. ABC was the most accurate for that image, with an error of less than 0.01. Nevertheless it produced erroneous estimates for other images, in one case underestimating  $\beta$  by 0.107 and overestimating for another image by 0.150. This simulation study has shown that all four algorithms are prone to getting stuck in low-probability regions of the parameter space.

## 5.2 Satellite remote sensing

For the satellite image described in Section 2.2 we used weakly informative priors  $\pi(\mu_j) \sim \mathcal{N}(\bar{y}, 5)$  and  $\pi(\sigma_j^2) \sim \mathcal{IG}(1, 0.01)$  for the  $k = 6$  mixture components and  $\pi(\beta) \sim \mathcal{U}(0, 3)$  for the inverse temperature. We were unable to evaluate accuracy for this image because the true values of the inverse temperature and the pixel labels are unknown. However, we can still compare the 95% HPD intervals for  $\beta$  obtained from the four algorithms, as well as the time taken to fit the model. These results are summarised in Table 2.

We ran pseudolikelihood and path sampling for a total of 60,000 iterations each, discarding the first 5,000 as burn-in. Due to the high computational cost of ABC-MCMC and the approximate exchange algorithm, we only ran these algorithms for 10,000 iterations with 5,000 burn-in, using 500 auxiliary iterations of Gibbs sampling. This took more than 50 hours, in comparison to 2 hours for pseudolikelihood and 1.4 hours for path sampling, plus 33 minutes to precompute the expectation of the sufficient statistic for 64 values of  $\beta$ . Pseudolikelihood was the fastest algorithm but it appears to have overestimated the value of  $\beta$  by a wide margin. There was no overlap between any of the credible intervals, but the other three algorithms all produced estimates in the neighbourhood of 1.3. This is higher than the critical value of  $\beta$ , which is 1.24 for a 2D image with  $k = 6$ . For comparison, we also analysed this image using PyMCMC [Strickland et al., 2012] and the R package `potts` [Geyer and Johnson, 2014]. PyMCMC produced a posterior estimate of 1.27, while the maximum

pseudolikelihood estimate (MPLE) returned by `potts` was 1.76.

### 5.3 Medical images

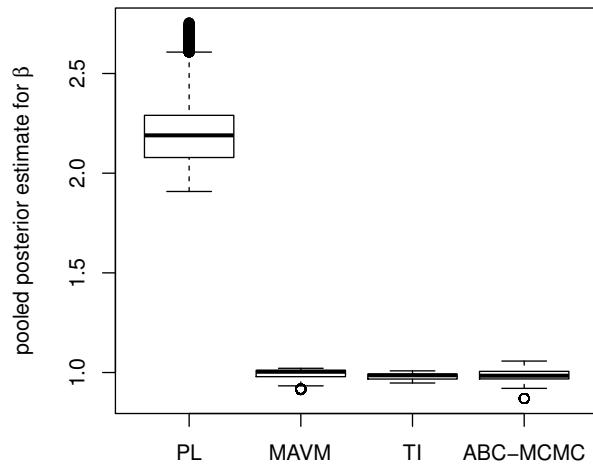
As explained in Section 2.3, the voxel intensity values in a CT scan are measured in Hounsfield units, with water at 0 HU and air at -1024 HU. We used informative priors for the noise parameters  $\pi(\mu_j) \sim \mathcal{N}(m_j, 200^2)$  with  $m_j \in \{-800, -600, -400, -200, 0, 200, 400, 600, 800\}$  and  $\pi(\sigma_j^2) \sim \mathcal{IG}(5, 5 \times 10^{-4})$  for the  $k = 9$  mixture components and  $\pi(\beta) \sim \mathcal{U}(0, 3)$  for the inverse temperature. The critical value of  $\beta$  for 3D images with  $k = 9$  is approximately 0.858, as shown in Figure 4b.

We ran pseudolikelihood and path sampling for 60,000 iterations on each of the 28 CT scans, but the approximate exchange algorithm and ABC were only run for 10,000, as we did for the satellite image and the synthetic data. The first 5,000 iterations were discarded as burn-in. The distributions of the MCMC samples for  $\beta$  are illustrated in Figure 9a. The pooled 95% HPD intervals were [2.207; 2.281] for pseudolikelihood, [0.934; 1.019] for the exchange algorithm, [0.988; 0.992] for path sampling, and [0.938; 1.034] for ABC. Pseudolikelihood has overestimated  $\beta$  in comparison to the other algorithms, as with the satellite image. The pooled credible intervals for the other three algorithms overlap with each other.

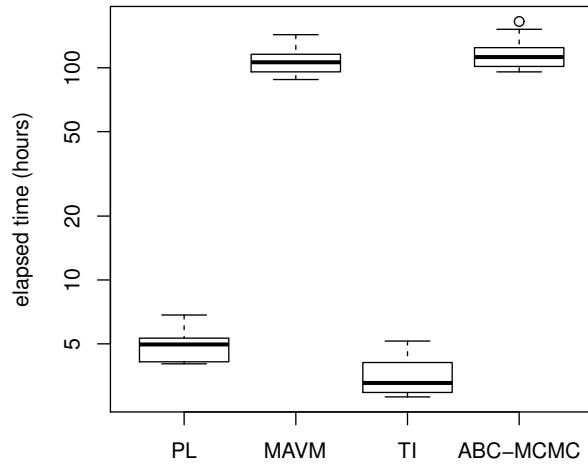
The distribution of runtimes are shown in Figure 9b. It took an average of 106.6 hours for 10,000 iterations of the exchange algorithm and 115 hours for ABC-MCMC. The average runtime for 60,000 iterations of pseudolikelihood was 4.9 hours, while the average for path sampling was 3.5 hours. It took 38 minutes to precompute 44 values of  $\mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$  for path sampling.

## 6 Concluding remarks

We have compared the approximate exchange algorithm, path sampling, pseudolikelihood, and ABC using a simulation study as well as real data. All of these Markov chain Monte Carlo algorithms were prone to becoming stuck in low-probability regions of the parameter space. Path sampling provided the best tradeoff between accuracy and computational cost over all of the images that we tested. This was due to precomputation of  $\mathbb{E}_{\mathbf{z}|\beta}[\mathbf{S}(\mathbf{z})]$  for fixed values of  $\beta$ , an operation that is highly parallelisable for modern computer architectures. The output of this precomputation step can be reused for multiple images with similar dimensions. The only drawback of path sampling is that it fails to account for the heteroskedasticity in the distribution of  $\mathbf{S}(\mathbf{z})$ . This can result in errors in the posterior distribution, particularly when  $\beta$  is above the critical value. [Moores et al. \[2015\]](#) introduced a precomputation step for ABC, with a second-order approximation that models the variance at the critical point. This produced a comparable runtime to path sampling, but without the outliers that are evident in Figure 8.



(a) Pooled posterior estimate for  $\beta$ .



(b) Elapsed time (hours).

Figure 9: Results for 28 CT scans of the ED phantom obtained using pseudo-likelihood (PL), the approximate exchange algorithm (MAVM), path sampling (TI), and approximate Bayesian computation (ABC-MCMC).

We have shown that simulating auxiliary variables at every iteration is too computationally expensive for applications in image analysis. Exact inference is infeasible for the scale of data that is commonly encountered in scientific studies, such as medical imaging and remote sensing. This has also been reported in previous studies, such as [McGrory et al. \[2009\]](#), [Everitt \[2012\]](#), [Moores et al. \[2015\]](#) and [Moores and Mengerson \[2014\]](#), but this is the first systematic comparison of the scalability of these algorithms. Methods such as [Prangle \[2014\]](#) can reduce the number of auxiliary iterations that need to be performed, but not enough to compensate for the two orders of magnitude difference in runtime that we observe.

Pseudolikelihood (PL) was the fastest of the four algorithms but produced very different estimates of  $\beta$  when applied to real data. The other three algorithms all depend on the value of the sufficient statistic  $S(\mathbf{z})$ , while PL uses the product of the conditional densities. We have illustrated how the approximation error of PL increases for values of  $\beta$  above the critical point. This is due to long-range dependence between the latent labels. Composite likelihood methods such as [Friel and Pettitt \[2004\]](#) and [Friel et al. \[2009\]](#) can reduce the approximation error by computing the estimate using sub-lattices. It is possible that the real images that we used in this study have a more complex correlation structure than what can be captured by the Potts model. This could explain the differences in the PL estimates between the simulation study and the real data.

This paper can serve as a guide to the selection of a suitable algorithm for Bayesian image analysis. We have provided our implementation of the four algorithms as an R package so that readers can experiment using their own data. There are faster algorithms for fitting a hidden Potts model, such as variational Bayes (VB) [[McGrory et al., 2009](#)] and iterated conditional modes (ICM) [[Besag, 1986](#)], but these provide limited options for estimating the inverse temperature. We have shown that it is feasible to use MCMC for images of realistic size, given reasonable computational power by contemporary standards. This enables the posterior distribution of the inverse temperature to be estimated along with all of the other model parameters.

## SUPPLEMENTARY MATERIAL

**R package bayesImageS:** R source package containing code to perform image segmentation using Markov chain Monte Carlo. Includes C++ implementations of all of the methods described in the article. (GNU zipped tar file)

**Simulation study:** Additional figures and tables of results for the simulation study of Section 5.1. (PDF)

## References

C. L. Alston, K. L. Mengersen, C. P. Robert, J. M. Thompson, P. J. Littlefield, D. Perry, and A. J. Ball. Bayesian mixture models in a longitudinal setting for analysing sheep CAT scan images. *Comput. Stat. Data Anal.*, 51:4282–4296, 2007. doi: 10.1016/j.csda.2006.05.013.

J. Besag. Spatial interaction and the statistical analysis of lattice systems. *J. R. Stat. Soc. Ser. B*, 36(2):192–236, 1974.

J. Besag. Statistical analysis of non-lattice data. *J. R. Stat. Soc. Ser. D*, 24(3):179–195, 1975.

J. Besag. On the statistical analysis of dirty pictures. *J. R. Stat. Soc. Ser. B*, 48(3):259–302, 1986.

E. O. Buzbas and N. A. Rosenberg. AABC: Approximate approximate Bayesian computation for inference in population-genetic models. *Theor. Popul. Biol.*, 99:31–42, 2015. doi: 10.1016/j.tpb.2014.09.002.

S. Cabras, M. Castellanos, and E. Ruli. A quasi likelihood approximation of posterior distributions for likelihood-intractable complex models. *METRON*, 72(2):153–167, 2014. doi: 10.1007/s40300-014-0040-5.

M.-H. Chen, Q.-M. Shao, and J. G. Ibrahim. *Monte Carlo Methods in Bayesian Computation*. Springer Series in Statistics. Springer-Verlag, New York, 2000.

S. R. Cook, A. Gelman, and D. B. Rubin. Validation of software for Bayesian models using posterior quantiles. *J. Comput. Graph. Stat.*, 15(3):675–692, 2006. doi: 10.1198/106186006X136976.

L. Cucala, J.-M. Marin, C. P. Robert, and D. M. Titterington. A Bayesian reassessment of nearest-neighbor classification. *J. Am. Stat. Assoc.*, 104(485):263–273, 2009. doi: 10.1198/jasa.2009.0125.

P. Del Moral, A. Doucet, and A. Jasra. An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Stat. Comput.*, 22(5):1009–20, 2012. doi: 10.1007/s11222-011-9271-y.

C. C. Drovandi and A. N. Pettitt. Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics*, 67(1):225–233, 2011. doi: 10.1111/j.1541-0420.2010.01410.x.

D. Eddelbuettel and C. Sanderson. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. *Comput. Stat. Data Anal.*, 71:1054–63, 2014. doi: 10.1016/j.csda.2013.02.005.

R. G. Everitt. Bayesian parameter estimation for latent Markov random fields and social networks. *J. Comput. Graph. Stat.*, 21(4):940–960, 2012. doi: 10.1080/10618600.2012.687493.

D. Feng. *Bayesian hidden Markov normal mixture models with application to MRI tissue classification*. PhD thesis, University of Iowa, 2008.

D. Feng and L. Tierney. *PottsUtils: Utility Functions of the Potts Models*, 2014. URL <http://CRAN.R-project.org/package=PottsUtils>. R package version 0.3-2.

N. Friel. Bayesian inference for Gibbs random fields using composite likelihoods. In C. Laroque, J. Himmelsbach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, editors, *Proc. Winter Simulation Conference*, pages 1–8, Dec 2012. doi: 10.1109/WSC.2012.6465236.

N. Friel and A. N. Pettitt. Likelihood estimation and inference for the autologistic model. *J. Comp. Graph. Stat.*, 13(1):232–246, 2004. doi: 10.1198/1061860043029.

N. Friel and A. N. Pettitt. Classification using distance nearest neighbours. *Stat. Comput.*, 21(3):431–437, 2011. doi: 10.1007/s11222-010-9179-y.

N. Friel, A. N. Pettitt, R. Reeves, and E. Wit. Bayesian inference in hidden Markov random fields for binary data defined on large lattices. *J. Comp. Graph. Stat.*, 18(2):243–261, 2009. doi: 10.1198/jcgs.2009.06148.

P. H. Garthwaite, Y. Fan, and S. A. Sisson. Adaptive optimal scaling of Metropolis-Hastings algorithms using the Robbins-Monro process. 2010. arXiv:1006.3690 [stat.ME].

A. Gelman and X.-L. Meng. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statist. Sci.*, 13(2):163–185, 1998. doi: 10.1214/ss/1028905934.

S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. PAMI*, 6:721–41, 1984.

C. J. Geyer and L. Johnson. *potts: Markov Chain Monte Carlo for Potts Models*, 2014. URL <http://CRAN.R-project.org/package=potts>. R package version 0.5-2.

P. J. Green and S. Richardson. Hidden Markov models and disease mapping. *J. Am. Stat. Assoc.*, 97:1055–1070, 2002.

A. Grelaud, C. P. Robert, J.-M. Marin, F. Rodolphe, and J.-F. Taly. ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–336, 2009. doi: 10.1214/09-BA412.

D. Hajduković. On the location of the critical point of the q-state Potts model on the hypercubic lattice. *J. Phys. A: Math. Gen.*, 16(6):L193, 1983. doi: 10.1088/0305-4470/16/6/004.

M. Huber. A bounding chain for Swendsen-Wang. *Random Struct. Algor.*, 22(1):43–59, 2003. doi: 10.1002/rsa.10071.

W. A. Kalender. *Computed Tomography: Fundamentals, System Technology, Image Quality, Applications*. Publicis, Erlangen, 3<sup>rd</sup> edition, 2011.

S. Z. Li. *Markov Random Field Modeling in Image Analysis*. Springer, Dordrecht, 3<sup>rd</sup> edition, 2009.

J.-M. Marin and C. P. Robert. *Bayesian core: a practical approach to computational Bayesian statistics*. Springer Texts in Statistics. Springer, New York, 2007.

J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian computational methods. *Stat. Comput.*, 22(6):1167–1180, 2012. doi: 10.1007/s11222-011-9288-2.

P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain Monte Carlo without likelihoods. *Proc. Natl Acad. Sci.*, 100(26):15324–15328, 2003. doi: 10.1073/pnas.0306899100.

C. A. McGrory, D. Titterington, R. Reeves, and A. N. Pettitt. Variational Bayes for estimating the parameters of a hidden Potts model. *Stat. Comput.*, 19(3):329–340, 2009. doi: 10.1007/s11222-008-9095-6.

C. A. McGrory, A. N. Pettitt, R. Reeves, M. Griffin, and M. Dwyer. Variational Bayes and the reduced dependence approximation for the autologistic model on an irregular grid with applications. *J. Comput. Graph. Stat.*, 21(3):781–796, 2012. doi: 10.1080/10618600.2012.632232.

P. Minvielle, A. Doucet, A. Marrs, and S. Maskell. A Bayesian approach to joint tracking and identification of geometric shapes in video sequences. *Image and Vision Computing*, 28(1):111–123, 2010. doi: 10.1016/j.imavis.2009.05.002.

J. Møller, A. N. Pettitt, R. Reeves, and K. K. Berthelsen. An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, 93(2):451–458, 2006. doi: 10.1093/biomet/93.2.451.

M. T. Moores and K. Mengersen. Bayesian approaches to spatial inference: modelling and computational challenges and solutions. In *Proc. 33<sup>rd</sup> Int. Wkshp MaxEnt*, volume 1636 of *AIP Conf. Proc.*, pages 112–117, Canberra, Australia, 2014. Amer. Inst. Physics. doi: 10.1063/1.4903701.

M. T. Moores and K. Mengersen. bayesImageS: Bayesian methods for image segmentation using a hidden Potts model, 2015. URL <https://researchdatafinder.qut.edu.au/display/n11957>. R package version 0.1-21.

M. T. Moores, C. C. Drovandi, K. Mengersen, and C. P. Robert. Pre-processing for approximate Bayesian computation in image analysis. *Stat. Comput.*, 25(1):23–33, 2015. doi: 10.1007/s11222-014-9525-6.

I. Murray, Z. Ghahramani, and D. J. C. MacKay. MCMC for doubly-intractable distributions. In *Proc. 22<sup>nd</sup> Conf. UAI*, pages 359–366, Arlington, VA, 2006. AUAI Press.

NASA. Landsat 7 science data users handbook. Technical report, National Aeronautics and Space Administration, 2011. URL <http://landsathandbook.gsfc.nasa.gov/>.

S. Okabayashi, L. Johnson, and C. J. Geyer. Extending pseudo-likelihood for Potts models. *Statistica Sinica*, 21:331–347, 2011.

A. N. Pettitt, N. Friel, and R. Reeves. Efficient calculation of the normalizing constant of the autologistic and related models on the cylinder and lattice. *J. R. Stat. Soc. Ser. B*, 65(1):235–246, 2003. doi: 10.1111/1467-9868.00383.

D. K. Pickard. Inference for discrete Markov fields: The simplest nontrivial case. *J. Am. Stat. Assoc.*, 82(397):90–96, 1987. doi: 10.1080/01621459.1987.10478394.

R. B. Potts. Some generalized order-disorder transformations. *Proc. Camb. Philos. Soc.*, 48:106–9, 1952. doi: 10.1017/S0305004100027419.

D. Prangle. Lazy ABC. *Stat. Comput.*, to appear, 2014. doi: 10.1007/s11222-014-9544-3.

J. G. Propp and D. B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algor.*, 9(1–2):223–252, 1996. doi: 10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.0.CO;2-O.

R. Reeves and A. N. Pettitt. Efficient recursions for general factorisable models. *Biometrika*, 91(3):751–757, 2004. doi: 10.1093/biomet/91.3.751.

T. Rydén and D. M. Titterington. Computational Bayesian analysis of hidden Markov models. *J. Comput. Graph. Stat.*, 7(2):194–211, 1998. doi: 10.1080/10618600.1998.10474770.

E. P. Simoncelli. Bayesian multi-scale differential optical flow. In B. Jähne, H. Haussecker, and P. Geissler, editors, *Handbook of computer vision and applications*, volume 2, chapter 14, pages 397–422. Academic Press, San Diego, 1999.

C. M. Strickland, R. J. Denham, C. L. Alston, and K. L. Mengersen. A Python package for Bayesian estimation using Markov chain Monte Carlo. In C. L. Alston, K. L. Mengersen, and A. N. Pettitt, editors, *Case Studies in Bayesian Statistical Modelling and Analysis*, Wiley Series in Probability and Statistics, chapter 25, pages 421–460. Wiley, 2012.

R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.*, 58:86–88, 1987. doi: 10.1103/PhysRevLett.58.86.

USGS. Landsat 8. Technical report, United States Geological Survey, Reston, VA, June 2014. URL <http://landsat.usgs.gov/landsat8.php>.

C. Varin, N. Reid, and D. Firth. An overview of composite likelihood methods. *Statistica Sinica*, 21:5–42, 2011.

R. D. Wilkinson. Accelerating ABC methods using Gaussian processes. In S. Kaski and J. Corander, editors, *Proc. 17th Int. Conf. AISTATS*, volume 33 of *JMLR W&CP*, pages 1015–1023, 2014.

G. Winkler. *Image Analysis, Random Fields and Markov chain Monte Carlo Methods: A Mathematical Introduction*. Springer-Verlag, Berlin Heidelberg, 2<sup>nd</sup> edition, 2003.